

Application Profiling on Cell-based Clusters

Hikmet Dursun,^{1,2} Kevin J. Barker,¹ Darren J. Kerbyson,¹ Scott Pakin¹

¹*Performance and Architecture Laboratory (PAL), Computer Science for HPC (CCS-1)
Los Alamos National Laboratory, NM 87545, USA*

²*Collaboratory for Advanced Computing and Simulations, Department of Computer Science
University of Southern California, CA 90089, USA
{hdursun, kjbarker, djk, pakin}@lanl.gov*

Abstract

In this paper, we present a methodology for profiling parallel applications executing on the IBM PowerXCell 8i (commonly referred to as the “Cell” processor). Specifically, we examine Cell-centric MPI programs on hybrid clusters containing multiple Opteron and Cell processors per node such as those used in the petascale Roadrunner system. Our implementation incurs less than 3.2 μ s of overhead per profile call while efficiently utilizing the limited local store of the Cell’s SPE cores. We demonstrate the use of our profiler on a cluster of hybrid nodes running a suite of scientific applications. Our analyses of inter-SPE communication (across the entire cluster) and function call patterns provide valuable information that can be used to optimize application performance.

1. Introduction

Application developers at the forefront of high-performance computing (HPC) have been investigating the use of hybrid architectures to improve application performance. Hybrid architectures attempt to improve application performance by combining conventional, general-purpose CPUs with any of a variety of more specialized processors such as GPUs, FPGAs, and Cells. The complexity stemming from hybrid architectures make understanding and reasoning about application performance difficult without appropriate tool support.

In this paper, we present a profiling library that can trace not only intra-Cell DMA events but also inter-Cell message passing. Our implementation is efficient in terms of resource consumption (only 12 KiB of SPE local store memory is required) and has an overhead of less than 3.2 μ s per profile call.

We employ a *reverse acceleration* programming model in which the hybrid cluster architecture is presented to the programmer as a logical cluster of Cell SPE processors by using the Cell Messaging Layer (CML) [1]. CML significantly reduces the effort needed to port applications to Cell clusters and has been used to port several scientific applications (e.g., the Sweep3D deterministic particle-transport kernel) to Los Alamos National Laboratory’s petascale Roadrunner supercomputer (comprising 6,120 dual-core Opterons plus 12,240 PowerXCell 8i processors). CML provides a subset of the functions and semantics of the MPI standard [2] including point-to-point communication, broadcasts, barriers, and global reductions.

The Cell processor’s complex architecture—eight synergistic processing elements (SPEs) managed by a single power processor element (PPE)—makes profiling tools essential for performance optimization. The IBM Cell Software Development Kit (SDK) [3] includes a Cell performance-debugging tool, PDT, that helps analyze the performance of a single Cell board (up to two Cell processors). PDT can trace only a specific set of SDK library functions such as SPE activation, direct memory access (DMA) transfers, synchronization, signaling, and user-defined events. Because PDT involves the slow PPE on the critical path of tracing, the PPE can easily become a performance bottleneck and may even influence application performance. Another tool for analyzing Cell performance is Vampir [4], which Nagel et al. used to visualize intra-Cell events such as mailbox communication and DMA transfers [5].

The key difference between our work and the works mentioned above is that we perform cluster-level analysis for MPI programs running on a hybrid architecture. The underlying message-passing model of CML, which treats an entire cluster of Opterons+Cells as a homogenous collection of SPEs,

has a central importance to our cluster-wide analysis. In addition to monitoring the same types of intra-Cell events as existing Cell profilers, our implementation can log inter-Cell, inter-blade, and inter-node communication. We have tested our implementation on up to 256 SPEs, although there is nothing limiting us from scaling up to thousands or even tens of thousands of SPEs.

Two parallel scientific applications—lattice Boltzmann (LB) flow simulation and atomistic molecular dynamics (MD) simulation—are used to test the profiler on the hybrid Opteron+Cell Roadrunner architecture using CML. Two sample uses of the profiler are also demonstrated: communication analysis and call-stack analysis.

The organization of the rest of this paper is as follows: Section 2 provides information about the Cell architecture, the Cell Messaging Layer, and our experimental testbed. Section 3 discusses the software design and implementation of our profiler software. Section 4 analyzes profiler performance using microbenchmarks and some sample applications. Finally, we summarize our study in Section 5.

2. Architectural background and testbed

In this section we describe the architecture of the PowerXCell 8i, the version of the Cell Broadband Engine that provides the bulk of the performance of our target cluster and the focus of our profiler study. We then briefly summarize the overall architecture of our testbed cluster. Finally, we describe the Cell Messaging Layer, which is an enabling technology for exploiting hybrid clusters and therefore a key insertion point for profiler events.

2.1. IBM PowerXCell 8i

The IBM PowerXCell 8i is the latest implementation of the Cell Broadband Engine. It drives the fastest supercomputer at the time of this writing, Roadrunner at Los Alamos [6], consists of a power processor element (PPE) and eight synergistic processing elements (SPEs) connected via an element interconnect bus (EIB), which supports a peak bandwidth of 204.8 GB/s for intra-chip data transfers among the PPE, SPEs, the memory, and the I/O interface controllers [7].

The PPE controls the eight SPEs, where each SPE contains a 3.2 GHz synergistic processing unit (SPU) core, 256 KB of a private, program-managed local store (LS) in place of a cache, and a memory flow controller (MFC) that provides DMA access to main memory. The SPE uses its LS for efficient instruction and data access, but it also has full access (via DMA)

to the coherent shared memory, including the memory-mapped I/O space.

To make efficient use of the EIB and to interleave computation and data transfer, the PPE and 8 SPEs are equipped with a DMA engine. Since an SPE’s load/store instructions can access only its private LS, each SPE depends exclusively on DMA operations to transfer data to and from the main memory and other SPEs’ local memories. The use of DMAs as a central means of intra-chip data transfer maximizes asynchrony and concurrency in data processing inside a Cell processor [8].

2.2. Testbed

Our testbed comprises 8 nodes, called *tri-blades*, where each tri-blade has two IBM QS22 Cell blades and one IBM LS21 AMD Opteron blade. The QS22 contains two PowerXCell 8i processors running at 3.2 GHz and each with an associated 4 GB of DDR2 memory. The LS21 blade includes two dual-core Opteron cores clocked at 1.8 GHz. Each tri-blade has a single connection to a Mellanox 4x DDR InfiniBand network. Typically, the Opterons handle mundane processing (e.g., file system I/O) while mathematically intensive elements are directed to the Cell processors. Each tri-blade in our testbed is architecturally identical to the tri-blades used in Roadrunner.

2.3. Cell Messaging Layer

CML is an implementation of common MPI functions for SPE-to-SPE communication in Cell-based clusters. The programming model underlying the CML is that applications run entirely on the SPEs. The SPE-centric model of CML assigns unique MPI ranks to each SPE assigned to an application. By means of using the PPE (and possibly conventional CPUs like Opterons if they exist in the cluster) primarily for transferring messages to SPEs in other blades instead of for computation, CML allows each SPE to communicate with other SPEs regardless of their location. In addition to handling communication operations, PPEs are also used to initialize CML, start programs on the SPEs and shut down CML after all SPEs invoke an *MPI_Finalize()*.

CML also provides Programmer’s Message Passing Interface (PMPI) functions [9] which have a one-to-one correspondence to MPI calls. This interface enables any calls made to the MPI functions, by the SPEs, to be intercepted and thus recorded. Section 3.2 discusses the use of PMPI calls within our profiler.

CML also offers a remote procedure call (RPC) mechanism through which SPEs can invoke a function on the PPE (PPEs can subsequently call a function on

the accompanying host CPU if it exists) and receive any results. This capability is particularly useful for our profiler, where local SPEs need to call a PPE *malloc()* to allocate space in PPE memory to hold the entire list of recorded events.

3. Software design details

Our implementation of the tracing library targets clusters of Cell processors. Each PPE within a Cell processor is responsible for synchronizing the program run on its SPEs. CML enables the total number of SPEs, as seen by an application, to scale: from a single processor containing eight SPEs to clusters of PlayStation3s [10], or to Roadrunner that contains 97,920 SPE cores. The remainder of this section outlines the design and implementation of the profiler including its memory use, and events that are profiled.

3.1. Data structures

The buffers that are used in the profiler implementation, along with the double-buffering operation of the buffers in the LS, is shown in Figure 1. This is discussed further in Section 3.2.




Figure 1. Operation of the profiler double-buffering implementation.

A cyclical pattern is used in Figure 1 to illustrate the allocation of buffers in LS. They switch roles repeatedly—while one is being used to record newly created events, the other is being dumped to PPE main memory. In comparison, the PPE memory layout is linear, where each small section, or event page, corresponds to the size of a single buffer in LS.

Table 1 summarizes the structure of profile events and of event pages that hold a number of events. It is crucial that the events and the buffers are allocated to fit the 16-byte boundary required for DMA

Table 1. Definitions of the data structures.

```
#define PAGE_SIZE 64
#define ALIGNED16 __attribute__((aligned(16)))

typedef enum { PROFILE_START,
    PROFILE_STOP, E, X, MPI_SEND, MPI_RECV,
    MPI_ALLREDUCE, MPI_REDUCE, MPI_BARRIER,
    MPI_BCAST, MFC_PUT, MFC_PUT64, MFC_GET64,
    MFC_PUT32, MFC_GET32} event_type_t;

typedef struct {
    double time_stamp;
    double duration;
    event_type_t type;
    unsigned long long enx;
    unsigned long long exx;
    short output_flag;
    int data[6] ALIGNED16;
} ALIGNED16 event_record_t;

typedef struct page_tag {
    struct page_tag* next_page;
    event_record_t events[PAGE_SIZE];
} ALIGNED16 event_page_t;
```

transactions. *ALIGNED16* is a short-hand notation for the *__attribute__((aligned(16)))* attribute, which specifies to the compiler to allocate the data structure to be 16 byte aligned. It is also important that 16 byte aligned profile data is structured the same both on SPE and PPE memories.

The enumerator lists the type of events our implementation can currently monitor. We record calls to the profiler start/stop functions, SPE function entry/exit (*E, X*), calls to the MPI functions implemented in CML (*MPI_SEND, MPI_RECV, MPI_ALLREDUCE, MPI_REDUCE, MPI_BARRIER, MPI_BCAST*) and various DMA put/get transactions which are issued by functions *spu_mfcdma32()* and *spu_mfcdma64()* defined in *libspe2*—the standard SPE library included in the IBM Cell/B.E. SDK. We have limited our implementation to cover only relevant DMA transaction types to our test applications.

In addition to recording the type of event, *event_record_t* also records a time stamp and the duration of an event, address of the called SPE function and its caller (*enx, exx*), an output flag to indicate that an event has happened and *data* array which includes destination/source, send/receive size and send/receive counts for MPI events. Effective addresses (*enx, exx*) are stored as an *unsigned long long* on both the SPE and PPE, so that they can be treated in a unified fashion no matter if the PPE code is compiled for 32-bit or 64-bit execution. One event record uses 80 bytes in memory.

A single buffer, or event page, is defined by *event_page_t*. The size of a page was set to be 64 in our testing (using 5,120 bytes). A pointer to the next

page to use is a part of *event_page_t* in case the current page fills up. Contrary to the SPE, which has two event pages, the PPE allocates a far greater number of event pages, as specified by *PAGE_PER_SPE*, for each SPE under its control. This was set to 10,000 giving a total PPE memory footprint of 400MiB (=8×10,000×64×80).

3.2. Implementation

CML based applications first start on the PPEs, which subsequently launches code on the SPEs. When the profiler is enabled, an instrumented SPE program, once launched, immediately invokes an allocation function on the PPE, using the CML's RPC mechanism, for event pages in main memory. Each SPE is returned the base address of the reserved memory via the same RPC mechanism. Before a SPE proceeds with actual application execution, it allocates two event buffers in its LS. However, this allocation is much smaller than its counterpart in main memory due to the limited size of the LS. In our tests the profiler statically allocates only two small event buffers of size 5,120 bytes, which holds up to 64 events, in SPE memory. Apart from the 10 KiB required for the two buffers, the profiler code requires less than an additional 2 KiB in LS but is dependent on the actual number of CML functions used by an application. This is ~30,000 times smaller than the memory use for event pages in the main memory of the PPE.

Profiler initialization is followed by the execution of the actual MPI application. Throughout the application run, the instrumented functions are called to record events. The instrumented operations, as provided by the profiler, create event logs. For instance, an SPE-to-SPE message-passing request invokes the corresponding instrumented MPI communication operation, which populates the event data structure with the relevant information, e.g., type, source/target, size of the message, and secondly calls the corresponding PMPI routine, which is implemented by CML, to send the actual message. The profiler library provides similar instrumented functions to profile other events including DMA operations and SPE function call activities.

SPE LS memory is limited to 256 KB. If it were to be filled with trace data, it would inhibit the execution of the SPE code. In order to circumvent this possibility, we use a double-buffering approach [11] to log trace events. Instead of continuously pushing events to a dynamically increasing allocation in LS, SPU writes profile event logs as they appear to one of the two small buffers allocated during profiler initialization. Once the buffer being used is full, previous buffer-dump operation is checked for

completion (Step 1 in Figure 1), by using *mfc_write_tag_mask* and *mfc_read_tag_status_al*, in order to avoid overwriting data being transferred. If the preceding dump has been completed, a non-blocking DMA (*mfc_put*) is issued to transfer the buffer to main memory (Step 2 in Figure 1). Each SPE sends the data to a privately reserved address, which it determines by using the memory base address received through the RPC mechanism during initialization, its local rank and number of previous dumps it has performed up to then. The SPE also switches the trace buffers and uses the available buffer to record new events (Step 3 in Figure 1). Meanwhile, the SPE execution continues without interruption as a non-blocking DMA is used. Once the second buffer is filled, the SPE switches buffers again and continues with recording events to the first buffer as it issues a DMA transfer (*mfc_put*) to dump the second buffer to the end of the preceding dump in the main memory (Step 4 in Figure 1). If the speed of event generation is faster than the time taken to transfer a single LS buffer to main memory then the application execution will pause. In such a case the size of the LS buffers can be increased but clearly at a reduction in the size of the LS store available to the application.


The double-buffering implementation not only overlaps data dumping with program execution, but also gives the capability of logging in excess of 10^4 times more events than the LS could have stored by using just two small buffers, and leaves more LS available for program and data in each SPE.

Upon the termination of tracing, the SPE program dumps the last buffer, regardless of how full it is, to main memory. Once all of the SPE terminate the PPE writes the profile data from main memory to several files, one per SPE, which contains the events which are ordered in terms of their time of occurrence. The output files can be post-processed for numerous performance analysis studies.

4. Results

Three applications were chosen to both quantify the overheads of the profiler use and also to illustrate its usefulness.

The first application is Sweep3D, which solves a single-group time-dependent discrete ordinates neutron-transport problem. It processes a regular three-dimensional data grid which is partitioned onto a logical two-dimensional processor array. Its computation consists of a succession of 3D wave fronts (sweeps), in which each processor receives boundary data from upstream neighbors, performs a computation on its local sub-grid, and produces



(a) Average cost per profiler call

(b) Slowdown due to the profiler (8 SPE run on a single Cell).

boundaries for downstream neighbors. All communications use MPI to transfer boundary data to neighboring processors.

The second application is a molecular-dynamics (MD) code [12]. The MD simulation follows the time evolution of the positions of N atoms by solving coupled ordinary differential equations. For parallelization, the MD code uses a 3-D spatial domain that is partitioned in all three dimensions into P subgrids of equal volume. Each step in the simulation requires the processing of the local sub-grid as well as boundary exchanges in each of 6 neighboring directions (i.e. the lower and higher neighbor subsystems in the x, y and z directions).

The third application is a lattice Boltzmann (LB) method for fluid flow simulations. The cellular-automata like application represents fluid by a density function on the grid points on a regular 3D lattice [10]. LB exhibits the same 3D communication pattern as for MD where each time step involves DF updates and inter-sub-grid density migrations.

4.1. Profiler performance overheads

The performance overhead of the profiler is dependent on the application as the mixture of communication and computation operations vary in each code. We therefore used an overhead metric by considering a worst-case by using a kernel application containing only communication calls and no computation. Additionally, by executing the kernel on a single Cell processor we ensure that only fast on-chip communications over the EIB are used. The

kernel application simply contained the communication pattern of the Sweep3D application thus resulting in a maximum rate of event generation.

An equal number of MPI send and receive calls, using a fixed size of 600 doubles (4,800 bytes), for the 8 SPEs on a single Cell was used for the results shown in Figure 2. Figure 2(a) shows the average overhead for each profiler call as a function of the number of events and Figure 2(b) shows the slowdown when varying the number of events. It can be seen that the average time required to record a single event is less than 3.2 μ s. This corresponds to a slowdown of a factor of 4.2 \times for large numbers of events as shown in Figure 2(b). Recall however that intra-cell communications take full advantage of the EIB which has a total bandwidth of 204.8 GB/s. Intra-cell communications using CML actually achieve a bandwidth of \sim 23 GB/s and latency of \sim 0.3 μ s as shown in Table 2. And hence a SPE-to-SPE message of size 4,800 bytes takes less than 1 μ s within a single Cell.

However, for a typical application running on a cluster of Cells, or on a hybrid processor configuration like Roadrunner, the SPE-to-SPE communications can be significantly more costly. For instance on Roadrunner, communications between SPEs on different nodes have a latency of over 11.7 μ s at small message bandwidth of 161 MB/s. Therefore, in practice, we expect the profiling overhead to be much lower due to the increasing cost of communications as well as computational activities in an application.




Figure 3. Communication pattern of Sweep3D.

4.2. Communication analysis

The information that is generated by the profiler is analyzed off-line. One log-file is generated for each SPE used by the application. Figure 3(a) shows an example of the SPE-to-SPE communication pattern of the original Sweep3D code on 8 nodes (32 SPEs per node) of a Roadrunner like cluster. In the figure, a larger square surrounded by thick lines and denoted by a node number, which contains 4×4 small squares, represents a tri-blade in the cluster. Each smaller square represents one Cell processor with 8 SPEs. The vertical and horizontal axes represent the sender and receiver SPE MPI ranks, and a colored pixel on the graph indicates a pair of communicating SPEs. The pixels are color coded to distinguish intra-node (blue) and inter-node (red) communications respectively.

The decomposition of Sweep3D's global grid onto a logical 2-D processor array can be seen in Figure 3(a). Each processor communicates with its neighbor in the logical x and y directions. For a 256 processor run, the 2-D processor array consists of 16×16 SPE processors. Each processor communicates with its x neighbors (± 1) as illustrated by the two sub-diagonals, and with its y neighbors (± 16) indicated by the outermost two off-diagonals. Message passing for the two x neighbors is performed on the same chip through high-bandwidth (25.6 GB/s) EIB; whereas communication with the y neighbors corresponds to some intra-node and some inter-node communications. These different inter-SPE communications incur different latency and bandwidth costs as shown in Table 2.

The high latency of inter-node communication in comparison to intra-cell communication stems from the involvement of PPes and Opterons in the former.

To achieve higher performance, parallel algorithms should be designed to exploit the low latency and high bandwidth of EIB connecting intra-cell SPEs and avoid inter-node communication wherever possible.

Figure 3(b) shows the communication pattern of a modified version of Sweep3D. In this implementation, much of the message passing is performed over the EIB. On each Cell, one SPE acts as a root and exclusively handles inter-node message passing by gathering messages from the other SPEs on the same chip and sending it to the root on the destination Cell. This reduces the number of inter-cell messages significantly and promises an increase in performance.

Figure 4 shows the communication pattern of MD for 256 SPE run. The logical arrangement of processors is in an $8 \times 8 \times 4$ processor array. Each SPE performs two intra-cell communications with x neighbors. Communications to y neighbors is comparably slower with half of the SPEs requiring inter-node communications. For example, in the first node, SPEs 1–8 and SPEs 25–32 have one of their y neighbors in the next node, while for SPEs 9–24 the communications to y neighbors only involves intra-node communications. For all SPEs, message passing with z neighbors is inter-node communication with a high communication cost. This suggests a possible optimization, to increase the number of communications over the EIB, as with Sweep3D.

Table 2. CML point-to-point performance.

Configuration	Latency	Bandwidth
Same Cell	0.272 μ s	22,994.2 MB/s
Same node	0.825 μ s	4,281.3 MB/s
Different nodes	11.771 μ s	161.2 MB/s




Figure 4. Communication pattern of MD.

As a matter of fact, event data structure as described in Section 3.1 has enough data to provide finer details on message passing events. For example inter-SPE and/or SPE-to-PPE communications can be analyzed in finer detail. Function use, duration, type of message passing activity, size of the message, type of data being sent (and/or received), count of a certain data type, and source/destination, can be analyzed to provide more insight into the program flow. It is also possible to automatically identify the communication pattern by measuring the degree of match between point-to-point communication matrix and predefined communication templates for regularly occurring communication patterns in scientific applications [13].

4.3. Call-stack analysis

The profiler library can also keep track of function entry and exits. This subsection illustrates the use of this functionality by a call-stack analysis as another use of our profiler.

Figure 5 shows the function call graph for the execution at the first SPE of a 256-SPE run for LB code. Instrumentation for 10 iterations is visualized and only a portion of call graph is provided for the clarity of presentation. The node shown as the root is the main function, which calls collision, streaming and communication functions once during every iteration. The nodes for these 3 functions include the source file name and the source code line information, which are looked up from a symbol table during post-processing. The node, which calls the *MPI_SEND/MPI_RECEIVE* implementation of CML, represent calls to the communication functions. Its children nodes show source/destination and data count of the message in parenthesis. The edges of the graph are marked with the number of times a particular event is observed. For instance, the node *MPI_SEND(0,56,132)* represents SPE 0 sends a message to SPE 56 of 132 bytes and it has occurred 10 times during the profiling.

The instrumentation is also done for functions expanded inline in other functions. The profiling calls indicate where the inline function is entered and exited. This requires that addressable versions of such functions must be available. A function may be given the attribute *no_instrument_function*, in which case this instrumentation will not be done. This can be used, for example, for high priority interrupt routines, and any function from which the profiling functions cannot safely be called, for example signal handlers.

The function call graph provides insight into program execution on a particular SPE on a cluster contributing to optimizations at the SPE level. In Figure 5, we have weighted the edges with function call numbers. Instead, operation completion time could be used as an alternative for weighting as the profiler keeps durations of events as well. A call graph can be used to identify bottlenecks of performance at the SPE level and shed a light on required algorithm modifications for improvement.




Figure 5. Function call graph for LB.

5. Conclusions

We have developed a low-memory-footprint (12 KiB of local store), minimally intrusive profiling library for parallel applications running on hybrid clusters of Opteron and Cell processors. Our library overlaps computations and DMA transfers to reduce application perturbation and efficiently utilizes the small amount of SPE local store available on Cell processors.

We have used our profiler library to analyze the performance of parallel scientific applications that run across multiple Cell processors, Cell blades, and cluster nodes. Inter-blade communication analysis for Sweep3D has shown how communication structure can affect application performance. We have ported two additional applications, LB and MD, to a hybrid Opteron+Cell cluster, and our profiler data suggests possible optimization opportunities. In order to demonstrate other uses of our library, we have analyzed the function-call pattern of a single SPE's program flow and used that to determine performance bottlenecks on the level of a SPE core.

While our study demonstrates high-speed, low-memory-overhead profiling for clusters augmented with Cell processors, it is certainly possible to optimize the profiler to further reduce its profiling cost and memory footprint. For example, the various types of profile events have different memory requirements (e.g., call-stack address records use only 16 bytes out of the 80 bytes allocated for the general event type). Therefore, restructuring the data types to be event-specific and adding compile-time options to customize the desired performance report may result in lower intrusion to program flow and reduce the post-processing effort for profile data.

Acknowledgements

This work was funded in part by the Advanced Simulation and Computing program of the Department of Energy. Los Alamos National Laboratory is operated by Los Alamos National Security LLC for the US Department of Energy under contract DE-AC52-06NA25396. We thank R. K. Kalia, P. Liu, A. Nakano, K. Nomura, R. Seymour, and P. Vashishta for providing us with their parallel MD and LB simulation codes.

References

- [1] S. Pakin, "Receiver-initiated Message Passing over RDMA Networks", in *Proceedings of the 22nd IEEE International Parallel and Distributed Processing Symposium (IPDPS 2008)*, Miami, Florida, April 2008.
- [2] M. Snir, S. Otto, S. Huss-Lederman, D. Walker, and J. Dongarra, *MPI: The Complete Reference*, volume 1, The MPI Core. The MIT Press, Cambridge, Massachusetts, 2nd edition, September 1998.
- [3] IBM. "Software Development Kit for Multicore Acceleration Version 3.0: Programmer's Guide".
- [4] H. Brunst and W. E. Nagel, "Scalable Performance Analysis of Parallel Systems: Concepts and Experiences", in *Proceedings of the Parallel Computing Conference (ParCo 2003)*, pages 737-744, Dresden, Germany, September 2003.
- [5] D. Hackenberg, H. Brunst, and W. E. Nagel, "Event Tracing and Visualization for Cell Broadband Engine Systems", in *Proceedings of 14th International Euro-Par Conference (Euro-Par 2008)*, pages 172-181, Las Palmas de Gran Canaria, Spain, August 2008.
- [6] K. J. Barker, K. Davis, A. Hoisie, D. J. Kerbyson, M. Lang, S. Pakin, and J. C. Sancho, "Entering the Petaflop Era: The Architecture and Performance and Roadrunner", in *Proceedings of IEEE/ACM SC08*, Austin, Texas, November 2008.
- [7] T. Chen, R. Raghavan, J. N. Dale, and E. Iwata, "Cell Broadband Engine Architecture and its First Implementation: A Performance View", *IBM Journal of Research and Development*, 51(5):559-572, September 2007.
- [8] M. Gschwind, H. P. Hofstee, B. Flachs, M. Hopkins, Y. Watanabe, and T. Yamazaki, "Synergistic Processing in Cell's Multicore Architecture", *IEEE Micro*, 26(2):10-24, March 2006.
- [9] S. Mintchev, and V. Getov, "PMPI: High-Level Message Passing in Fortran77 and C", in *Proceedings of the International Conference and Exhibition on High-Performance Computing and Networking (HPCN 1997)*, volume 1225 of *Lecture Notes in Computer Science*, pages 603-614, 1997. Springer.
- [10] K. Nomura, S. W. de Leeuw, R. K. Kalia, A. Nakano, L. Peng, R. Seymour, L. H. Yang, and P. Vashishta, "Parallel Lattice Boltzmann Flow Simulation on a Low-cost Playstation 3 Cluster", *International Journal of Computer Science*, 2008.
- [11] J. C. Sancho, and D. J. Kerbyson, "Analysis of Double Buffering on two Different Multicore Architectures: Quad-core Opteron and the Cell-BE", in *Proceedings of the 22nd IEEE International Parallel and Distributed Processing Symposium (IPDPS 2008)*, Miami, Florida, April 2008.
- [12] A. Nakano, R. K. Kalia, K. Nomura, A. Sharma, P. Vashishta, F. Shimojo, A. C. T. van Duin, W. A. Goddard, R. Biswas, D. Srivastava, and L. H. Yang, "De Novo Ultrascale Atomistic Simulations On High-End Parallel Supercomputers", *International Journal of High Performance Computing Applications*, 22(1):113-128, 2008.
- [13] D. J. Kerbyson, K. J. Barker, "Automatic Identification of Application Communication Patterns via Template", *18th International Conference on Parallel and Distributed Computing Systems (ISCA PDCS 2005)*, pages 114-121, Las Vegas, Nevada, September 2005.